Chapter 8

The theory of the thermodynamic limit

8.1 Introduction

As far as I know, Khinchin's programme of using the central limit theorem in order to derive asymptotic expressions for the probability measure has not been pursued further. Instead, the attention has shifted to the so-called theory of the thermodynamic limit (TTL), which also exploits the fact that one deals in statistical mechanics with large numbers of particles. This is a far more powerful theory due to its applicability to very general types of interaction and to phase transitions.

The first proof of the existence of the thermodynamic limit was given by Van Hove in 1949, but the theory has been developed mainly in the 1960s and 1970s. Most of the results are collected in Ruelle's book 'Statistical Mechanics: Rigorous Results' (Ruelle 1969). A survey which is less complete but much more accessible is (Lanford 1973), on which much of the following is based.

The "thermodynamic limit" is the limit in which the number of particles and the volume go to infinity, while their ratio is or approaches a constant, finite value, and also the energy per particle is constant. In this limit, one typically studies the behaviour of the probability that a certain phase function, for instance the energy, has a value in a given interval. What one wants to show is that this probability becomes concentrated on one value of the phase function. Under some suitable restrictions on the phase function, this result can indeed be obtained. One can also show that the probability that the phase function has some other value falls off exponentially with the number of particles. It is also possible that the probability becomes concentrated on a range of values, instead of on a single point. This situation corresponds to a phase transition.

The importance of such results is clear. Generally the value of a phase function can vary wildly across phase space. Thus, outcomes of measurements of phase functions corresponding to a thermodynamic observable may generally vary. But in the presence of the mentioned results of the TTL, outcomes will be found near one value. Clearly this is helpful to understand why the phase averaging method works (note that if a phase function has its value in a certain interval, also its phase average will lie in that interval). Another way to look at these results is the following. In general, many variables are necessary in order to characterise a system consisting of a large number of particles. However, in thermodynamics only a few of them are already sufficient. The same now turns out to be true in statistical physics, after taking the thermodynamic limit, since (as will turn out) the particle density and the energy per particle alone are sufficient to determine the characteristics of the system.

The connection of the TTL with Khinchin's theory is obvious. In both cases the large number of particles plays a key role. The difference lies in the limit theorems that are employed. In the technical details there are of course more differences; for instance, in the limiting process not only the number of particles, but also the volume tends to infinity. The manner in which the volume is increased is of importance here, although the ultimate result turns out (fortunately) to be rather insensitive to it.

A second correspondence is that the result is valid for special phase functions only. We have seen that Khinchin argued that this is an advantage, since general ergodic theory is too general, and we need to have ergodic results only for functions that are actually encountered in statistical mechanics. In his case, the restricted class of functions are the sum functions, which is too limited; here, other restrictions are applied, the most important of which are stability and finite range. The restrictions are inspired by physical considerations. Moreover, the class of functions is much wider than just sum functions. Therefore, the theory is applicable to a broader class of physically interesting systems.

This leads to an important difference with Khinchin's programme, namely that there is no methodological paradox. We have seen previously that the restriction to sum functions leads, in case of the Hamiltonian, to the conclusion that strictly speaking no interaction between the molecules is allowed for. In the present case this restriction is lifted, and many interesting types of interaction are included in the theory.

A second difference is that there is no equivalent of Khinchin's ergodic theorem; in fact, time averages now play no role at all. The justification of why microcanonical averages may be used as analogues to thermodynamical quantities follows a different scheme. Instead of demonstrating the equivalence of ensemble and time averages, the object is now to show that the values of observables are to a high degree insensitive to the exact microstate. This would imply that considering not phase functions but their ensemble averages hardly ever leads to large errors.

Thus, we can say that the TTL has its own answer to our fundamental question of why the use of ensemble averages works. In comparison to the orthodox ergodic theory, it is far better off. The restriction to metrically transitive systems has been removed; only the measure zero problem remains. The theory is in many respects a successor of the Khinchin programme, even though no ergodic theorem is used.

8.2 Lanford's theory of the thermodynamic limit

Let us now look in more detail at the theorems that can be derived and the suppositions that go into them. I will follow the article (Lanford 1973) (especially the introductory section), since this is a clear account of the theory, without too much complicating details which are not necessary for the present purposes. In other sources (for instance (Ruelle 1969)) slightly different theorems may be found: Both the conditions on the observables and the sense in which one lets the volume go to infinity may be a bit different from Lanford's theorems.

Whereas Khinchin studies an asymptotic expression for the microcanonical probability distribution and only later makes use of special properties of functions of interest in physics, Lanford explicitly studies the distribution of phase functions that obey certain requirements inspired by the application to statistical mechanical systems from the start. Specifically, he studies the asymptotic behaviour of

$$\mu(\{x : F(x)/N \in J\}), \tag{8.1}$$

i.e. the microcanonical probability that a certain phase function F/N has values in a certain region J. The division by the particle number N is because the phase functions in question are supposed to scale with N. Here $x = (q_1, \dots, q_N)$ is a point in configuration space, not phase space; this is not a fundamental restriction, but it is introduced for convenience only. F may be a vector valued function.

As noted before, the observables for which Lanford proves the existence of the thermodynamic limit are of a special type. They should obey the list of conditions below. Furthermore, it is essential in this approach that the potential energy U is among the observables that satisfy these conditions.

a) Continuity	For each N, $F(q_1,\ldots,q_N)$ is a continuous function
b) Symmetry	For each N, $F(q_1,\ldots,q_N)$ is a symmetric function
c) Translation invariance	For each N and each $a \in \mathbb{R}^3$, $F(q_1+a,\ldots,q_N+a) = F(q_1,\ldots,q_N)$
d) Normalisation	$F(q_1) = 0$
e) Finite range	There exists a real number $R < \infty$ such that, if $\|q_i - q'_j\| > R \;\forall i,j$ then $F(q_1,\ldots,q_N,q'_1,\ldots,q'_M) = F(q_1,\ldots,q_N) + F(q'_1,\ldots,q'_M)$

In case F is vector valued, the last condition should be modified, in the sense that the range R of F is then defined to be the supremum of the ranges of the components of F.

Several comments are in order. Some of these conditions (continuity, translation invariance and normalisation) are rather harmless. The condition of symmetry is very powerful; it allows one to exploit the fact that systems of a large number of indistinguishable particles are considered. We have seen this before in the discussion of equilibrium in Boltzmann's approach to SM, where it was proved on the basis of symmetry arguments that one single macrostate occupies an overwhelmingly large part of phase space, because this macrostate can be realised in a huge number of ways by permutations of the particles.

Also the condition of finite range deserves special attention. The physical motivation (when F represents the potential energy) is that parts of the system that are sufficiently far away from each other, do not interact. In other versions of the TTL (see Ruelle), this condition is weakened a bit, by allowing decreasing but nonzero interaction for large distances between the parts. Further, note that the formulation of the condition of finite range resembles that of extensivity (roughly, a quantity is extensive if it scales with the size of the system and intensive if it is independent of the size of the system). Thus, for intensive quantities the condition of finite range is false.

In the following, we will suppose that the potential energy U satisfies the above conditions. We will also consider another (possibly vector valued) observable F, which satisfies these conditions as well. Our goal is now to study the (Lebesgue) probability that both U/N and F/N have values within given ranges, and do this for increasing particle number N. The range of values of U/N (denoted by I_ε) is taken to be a narrow interval of values around $U/N = \varepsilon$.

So, the function that we are interested in is

$$V(\Lambda,N,U,I_\varepsilon,F,J) \quad = \quad \frac{1}{N!}\mu\left(\left\{(q_1,\ldots,q_N) \in \Lambda^N : \frac{U(q_1,\ldots,q_N)}{N} \in I_\varepsilon\right.\right.$$

$$\text{and } \frac{F(q_1,\ldots,q_N)}{N} \in J \Big\}\Big) .$$

Here Λ denotes the region in three dimensional space to which the positions of the particles are confined; the volume of this region will be denoted by $V(\Lambda)$. The function V thus gives the Lebesgue measure (divided by $N!$) of the set of points for which both the energy U and the phase function F take values in specified sets. In the following, we will study the asymptotic properties of this function.

The phrase "the existence of the thermodynamic limit" refers to the limit of the logarithm of this function, divided by the number of particles. If it exists, it is denoted by

$$s(\rho,U,I_\varepsilon,F,J) = \lim \frac{1}{N} \log V(\Lambda,N,U,I_\varepsilon,F,J), \tag{8.2}$$

which only depends on the particle number and the volume through their ratio $\rho = \lim N/V(\Lambda)$.

Of course, the exact form of the limiting process is of importance here, especially the sense in which the volume goes to infinity. Lanford lets the volume go to infinity in the sense of Van Hove, and assumes that it is approximable by rectangles. Roughly speaking, this means that the volume of the boundary will become negligible w.r.t. the volume of the interior (for a more precise description, see (Lanford 1973, pp. 28–33)). As noted above, the thermodynamic limit also involves the particle density approaching a constant value, and the energy per particle being nearly constant. Lanford then proves that either the limit (8.2) exists and is finite, or V goes to zero faster than exponentially, corresponding to $s = -\infty$. The latter case need not worry us, since it only means that for this particular choice of J the probability that $F/N \in J$ goes to zero very quickly.

The function $s(\rho,I_\varepsilon,U,F,J)$ is a function of open sets J. Next (suppressing the dependence on I_ε, U, ρ and F in the notation), we define a related function which depends on the values x:

$$\tilde{s}(x) = \inf\{s(J) : \text{ open sets } J \ni x\} \tag{8.3}$$

Then, interestingly, a converse relation can be proved:

$$s(J) = \sup\{\tilde{s}(x) : x \in J\}. \tag{8.4}$$

Thus, only those points x for which $\tilde{s}(x)$ is maximal contribute to $s(J)$.

Suppose now, that we compare two intervals J_1 and J_2. From (8.2) we see that, asymptotically, the probability that $F \in J_1$ and that $F \in J_2$ are in the proportion of

$$\frac{\mu(F/N \in J_1)}{\mu(F/N \in J_2)} = \frac{V(J_1)}{V(J_2)} \overset{N\to\infty}{\approx} e^{N(s(J_1)-s(J_2))}. \tag{8.5}$$

If $s(J_1)$ is strictly smaller than $s(J_2)$, this goes to zero exponentially in N. From the fact that $s(J)$ can be written as a supremum (8.4), we see that all probability mass gets concentrated on the set of points where $\tilde{s}(x)$ is maximal.

Of the function $\tilde{s}(x)$ two important properties can be demonstrated, which together give very interesting information about the form of $\tilde{s}(x)$ and therefore about the distribution of F in the thermodynamic limit. First, $\tilde{s}(x)$ is a concave function of x. Secondly, it is upper semi-continuous. Together, these properties assure that $\tilde{s}(x)$ can have at most one supremum, which it attains either nowhere, in a single point, or in an interval. Thus we see that the probability that the observable F has a value outside the point or interval where $\tilde{s}(x)$ is maximal, will tend to zero in the thermodynamic limit.

Let us go through the three possible cases. First, $\tilde{s}(x)$ may attain its supremum in a single point $x = x_0$. The probability that the phase function F has a value larger than $x_0 + \varepsilon$ for any $\varepsilon > 0$ is then given by

$$\begin{aligned} \mu(F/N \in (x_0+\varepsilon,\infty)) &\overset{N\to\infty}{\approx} e^{N(s((x_0+\varepsilon,\infty))-s((-\infty,\infty)))} \\ &= e^{N(\tilde{s}(x_0+\varepsilon)-\tilde{s}(x_0))}, \end{aligned} \tag{8.6}$$

where $s((-\infty,\infty))$ in the exponent accounts for the normalisation of the probability, and (8.4) has been used in the second line. Since $\tilde{s}(x)$ has its single maximum in $x = x_0$, this probability goes to zero exponentially in N. Similarly, the probability that F/N is smaller than $x_0 - \varepsilon$ goes to zero.

Secondly, $\tilde{s}(x)$ may attain its supremum in an interval (x_0, x_1). Then, a similar argument shows that the probabilities that F/N is smaller than $x_0 - \varepsilon$, or larger than $x_1 + \varepsilon$, go to zero exponentially in N. However, nothing general can be said about the distribution within the interval. This case is said to correspond to a phase transition.

Third, $\tilde{s}(x)$ may never take on its supremum. In this case the probability that $F < x_0$ will, for any x_0, go to zero. So the probability mass "escapes to infinity", and also the mean value of F/N diverges. It can be shown (Ruelle 1970) that this, unwanted, behaviour is ruled out when the interaction $U(q_1,\ldots,q_N)$ satisfies a condition called superstability.

In conclusion, we have seen that for the Lebesgue measure on a shell of nearly constant potential energy, under the specified conditions (including superstability) on the observables, all probability will in the thermodynamic limit concentrate on a single value of F/N (which Lanford calls the equilibrium value of the observable), or on an interval of values in case of a phase transition. This also means that any phase function obeying the conditions will, in the absence of phase transitions, equal its ensemble average with a probability approaching one. This fact can be used in the justification of the use of phase averages, which was the starting point of our discussion.

8.3 Discussion

What does an answer to the Explanandum on the basis of the theory of the thermodynamic limit look like? In contrast to the previous accounts measurement outcomes are now taken to correspond with instantaneous values of phase functions instead of their time averages. The TTL says that these phase functions, if they obey certain conditions motivated by physics, and if volume and particle number are large, will with high probability have values near their equilibrium value. Thus the values of the observables will in the limit be insensitive to the precise microstate of the system. Since in such a case also the microcanonical phase average will be the same as (or very close to) the equilibrium value, this explains why these phase averages work.

In my view the importance of the TTL in the discussion of why the averaging method works has been underestimated. For instance, Sklar writes:

> 'Let us first note that the task here [in the TTL] is not to justify the equilibrium ensemble distribution. This is presupposed.' (Sklar 1993, p. 80)

and consequently does not pay much attention to the TTL in his survey of the foundations of statistical mechanics. Indeed, a justification of the microcanonical measure is not what the authors in this field set as their goal. But this does not mean that their results cannot be used for that purpose. In fact, we have just seen an explanation of why averages in the Lebesgue measure can be used as a description of a single system. I submit that the TTL even gives a better answer to the Explanandum than the accounts discussed in the previous chapters, the main improvement in comparison with Khinchin being the necessary dynamical assumptions, which are both more general and more realistic.

But of course the TTL does not give a complete, non-circular justification of the Lebesgue measure. Lanford is very clear about that:

> '...it becomes a purely mathematical problem (and perhaps not a totally hopeless one) to prove that with very high probability any observable is near its equilibrium value, probability always being computed with respect to Lebesgue measure in configuration (or phase) space.

> It is a much more profound problem to understand why events which are very improbable with respect to Lebesgue measure do not occur in nature. I, unfortunately, will have nothing to say about this latter problem.' (Lanford 1973, p. 2)

Thus he acknowledges the measure zero problem, and admits that the TTL cannot solve it. But my point is that it is not worse off than the other approaches we encountered so far.

Indeed, with respect to the measure zero problem the same comments that were made in connection with Khinchin's theory apply here also (see page 88). First, for finite systems of any size the set of points for which the observables have a value different from their equilibrium value still has finite, although small, measure. Therefore we are dealing again with the "measure epsilon problem". That is, what we are looking for are reasons why small Lebesgue measure corresponds with small probability of occurrence in nature. Secondly, since the TTL is meant to apply to systems that need not be metrically transitive, Farquhar's objection applies here also (see the discussion on page 88). However, in the context of the TTL this problem may be remedied more easily. Rather than the Lebesgue measure on a thickened energy shell, one could consider the Lebesgue measure on that part of phase space that is singled out by specifying the values of all invariants of the motion. Applying the TTL to the distribution of observables with respect to this restricted measure will meet Farquhar's objection.

Turning now to the list of ingredients of an ideal explanation, the important thing to notice is that the dynamical assumptions that are used in the TTL are very different from both the ergodic approach and Khinchin's approach. No use is made of metrical transitivity or other ergodic properties. Also, there is no restriction to sum functions, and consequently no methodological paradox. Instead the observables for which the existence of the thermodynamic limit is proved have to obey the assumptions listed in section 8.2. These assumptions are reasonable for the interaction energy and for other extensive observables.

Of course the large number of particles has been used explicitly in this approach. The fact that the Explanandum refers to equilibrium, however, has not been used, and this is a severe shortcoming. Finally, with respect to isolation again the same comments apply as in the two previous cases: This has been taken into account by considering an energy surface (or shell), and time-independent Hamiltonian.

There is yet another difference with the ergodic approach. In the ergodic theorems, some special property of the microcanonical measure is demonstrated, since it is the unique (stationary and absolutely continuous) measure which yields averages that coincide with time averages. In this limited sense the microcanonical measure is *derived*. In Lanford's theory on the other hand, the Lebesgue measure needn't be the unique measure that gives measure one to sets where $\tilde{s}(x)$ is maximal. In fact, it is easily seen that it is not, since the canonical measure has a similar property. Therefore, one has even less reason to conclude that the microcanonical measure is the correct, real measure. But even if the microcanonical measure is not unique in this respect, one may be able to explain the success of its use.

Of course, for such an explanation to be fully satisfactory the measure zero problem has to be solved. Moreover, it is clear that the fact that the system is in equilibrium has to be used in a solution of this problem, because so far in the TTL no reference to equilibrium has been made. In the next chapter we will encounter a completion of the TTL with an argument aiming to show that sets of small microcanonical measure may be neglected in practice, if the system is in phenomenological equilibrium.